arXiv:1510.03945v1 [cs.DS] 14 Oct 2015

An $O(\log \mathrm{OPT})$-approximation for covering and packing minor models of θ_r*

Dimitris Chatzidimitriou† Jean-Florent Raymond‡§¶ Ignasi Sau‡
Dimitrios M. Thilikos†‡

Abstract

Given two graphs G and H, we define $\mathsf{v\text{-}cover}_H(G)$ (resp. $\mathsf{e\text{-}cover}_H(G)$) as the minimum number of vertices (resp. edges) whose removal from G produces a graph without any minor isomorphic to H. Also $\mathsf{v\text{-}pack}_H(G)$ (resp. $\mathsf{e\text{-}pack}_H(G)$) is the maximum number of vertex- (resp. edge-) disjoint subgraphs of G that contain a minor isomorphic to H. We denote by θ_r the graph with two vertices and r parallel edges between them. When $H = \theta_r$, the parameters $\mathsf{v\text{-}cover}_H$, $\mathsf{e\text{-}cover}_H$, $\mathsf{v\text{-}pack}_H$, and $\mathsf{e\text{-}pack}_H$ are NP-hard to compute (for sufficiently big values of r). Drawing upon combinatorial results in [CRST15], we give an algorithmic proof that if $\mathsf{v\text{-}pack}_{\theta_r}(G) \leq k$, then $\mathsf{v\text{-}cover}_{\theta_r}(G) = O(k \log k)$, and similarly for $\mathsf{e\text{-}pack}_{\theta_r}$ and $\mathsf{e\text{-}cover}_{\theta_r}$. In other words, the class of graphs containing θ_r as a minor has the vertex/edge Erdős-Pósa property, for every positive integer r. Using the algorithmic machinery of our proofs we introduce a unified approach for the design of an $O(\log \mathrm{OPT})$-approximation algorithm for $\mathsf{v\text{-}pack}_{\theta_r}$, $\mathsf{v\text{-}cover}_{\theta_r}$, $\mathsf{e\text{-}pack}_{\theta_r}$, and $\mathsf{e\text{-}cover}_{\theta_r}$ that runs in $O(n \cdot \log(n) \cdot m)$ steps. Also, we derive several new Erdős-Pósa-type results from the techniques that we introduce.

Keywords: Erdős-Pósa property, packings in graphs, coverings in graphs, approximation algorithms, protrusion decomposition.

2000 MSC: 05C70.

*The second author has been partially supported by the Warsaw Centre of Mathematics and Computer Science and by the (Polish) National Science Centre grant PRELUDIUM 2013/11/N/ST6/02706. The first and the last authors were co-financed by the E.U. (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: "Thales. Investing in knowledge society through the European Social Fund". Email addresses: hatzisdimitris@gmail.com, jean-florent.raymond@mimuw.edu.pl, ignasi.sau@lirmm.fr, and sedthilk@thilikos.info.

†Department of Mathematics, National and Kapodistrian University of Athens, Athens, Greece.
‡AlGCo project team, CNRS, LIRMM, France.
§University of Montpellier, Montpellier, France.
¶Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Warsaw, Poland.

1 Introduction

All graphs in this paper are undirected, do not have loops but they may contain multiple edges. We denote by θ_r the graph containing two vertices x and y connected by r parallel edges between x and y. Given a graph class $\mathcal{C}$ and a graph G, we call $\mathcal{C}$-*subgraph* of G any subgraph of G that is isomorphic to some graph in $\mathcal{C}$. All along this paper, when giving the running time of an algorithm with input some graph G, we agree that $n = |V(G)|$ and $m = |E(G)|$.

Packings and coverings. Paul Erdős and Lajos Pósa, proved in 1965 [EP65] that there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that for each positive integer k, every graph either contains k vertex-disjoint cycles or it contains $f(k)$ vertices that intersect every cycle in G. Moreover, they proved that the "gap" of this min-max relation is $f(k) = O(k \cdot \log k)$ and that this gap is optimal. This result initiated an interesting line of research on the duality between coverings and packings of combinatorial objects. To formulate this duality, given a class $\mathcal{C}$ of connected graphs, we define $\mathsf{v\text{-}cover}_{\mathcal{C}}(G)$ (resp. $\mathsf{e\text{-}cover}_{\mathcal{C}}(G)$) as the minumun cardinality of a set S of vertices (resp. edges) such that each $\mathcal{C}$-subgraph of G contains some element of S. Also, we define $\mathsf{v\text{-}pack}_{\mathcal{C}}(G)$ (resp. $\mathsf{e\text{-}pack}_{\mathcal{C}}(G)$) as the maximum number of vertex- (resp. edge-) disjoint $\mathcal{C}$-subgraphs of G.

We say that $\mathcal{C}$ has the *vertex Erdős–Pósa property* (resp. the *edge Erdős–Pósa property*) if there is a function $f : \mathbb{N} \to \mathbb{N}$, called *gap function*, such that, for every graph G, $\mathsf{v\text{-}cover}_{\mathcal{C}}(G) \leq f(\mathsf{v\text{-}pack}_{\mathcal{C}}(G))$ (resp. $\mathsf{e\text{-}cover}_{\mathcal{C}}(G) \leq f(\mathsf{e\text{-}pack}_{\mathcal{C}}(G))$). Using this terminology, the original result of Erdős and Pósa says that the set of all cycles has the vertex Erdős–Pósa property with gap $O(k \cdot \log k)$. The general question in this branch of Graph Theory is to detect instantiations of $\mathcal{C}$ which have the vertex/edge Erdős–Pósa property (in short, $\mathsf{v/e}$-*EP-property*) and when this is the case, minimize the gap function f. Several theorems of this type have been proved concerning different instantiations of $\mathcal{C}$ such as odd cycles [KN07, RR01], long cycles [BBR07], and graphs containing cliques as minors [DiKW12] (see also [RRST96, KKK12, GK09] for results on more general combinatorial structures).

A general class that is known to have the v-EP-property is the class $\mathcal{C}_H$ of the graphs that contain some fixed planar graph H as a minor[1]. This fact was proven by Robertson and Seymour in [RS86] and the best known general gap is $f(k) = O(k \cdot \log^{O(1)} k)$ due to the results of [CC13] — see also [FJW13, FJS13] for better gaps for particular instantiations of H. Moreover, the planarity of H appears to be the right dichotomy, as for non-planar H's, $\mathcal{C}_H$ *does not* satisfy the v-EP-property. Besides the near-optimality of the general upper bound of [CC13], it is open whether the lower bound $\Omega(k \cdot \log k)$ can be matched for the general gap function, while this is indeed the case when $H = \theta_r$ [FJS13].

The question about classes that have the e-EP-property has also attracted some

[1] A graph H is a *minor* of a graph G if it can be obtained from some subgraph of G by contracting edges.

attention (see [BBR07]). According to [Die05, Exercice 23 of Chapter 7], the original proof of Erdős and Pósa implies that cycles have the e-EP-property with gap $O(k \cdot \log k)$. Moreover, as proved in [RST13], the class $\mathcal{C}_{\theta_r}$ has the e-EP-property with the (non-optimal) gap $f(k) = O(k^2 \cdot \log^{O(1)} k)$. Interestingly, not much more is known on the graphs H for which $\mathcal{C}_H$ has the e-EP-property and is tempting to conjecture that the planarity of H provides again the right dichotomy. Other graph classes that are known to have the e-EP-property are rooted cycles [PW12] (here the cycles to be covered and packed are required to intersect some particular set of terminals of G) and odd cycles for the case where G is a 4-edge connected graph [KK12], a planar graph [KV04], or a graph embeddable in an orientable surface [KN07].

Approximation algorithms. The above defined four graph parameters are already quite general when $\mathcal{C} := \mathcal{C}_H$. From the algorithmic point of view, the computation of $\mathsf{x\text{-}pack}_{\mathcal{C}_H}$ (for $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$) corresponds to the general family of graph packing problems, while the computation of $\mathsf{x\text{-}cover}_{\mathcal{C}_H}$ belongs to the general family of graph modification problems where the modification operation is the removal of vertices/edges (depending on whether $\mathsf{x} = \mathsf{v}$ or $\mathsf{x} = \mathsf{e}$). Interestingly, particular instantiations of $H = \theta_r$ generate known, well studied, NP-hard problems. For instance, asking whether $\mathsf{v\text{-}cover}_{\mathcal{C}_{\theta_r}} \leq k$ generates Vertex Cover for $r = 1$, Feedback Vertex Set for $r = 2$, and Diamond Hitting Set for $r = 3$ [FJP10, FLMS12]. Moreover, asking whether $\mathsf{x\text{-}pack}_{\mathcal{C}_{\theta_r}}(G) \geq k$ corresponds to Vertex Cycle Packing [BTY11, KLL02] and Edge Cycle Packing [ACR03, KNS+07] when $\mathsf{x} = \mathsf{v}$ and $\mathsf{x} = \mathsf{e}$, respectively. Finally, asking whether $|E(G)| - \mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_r}}(G) \leq k$ corresponds to the Maximum Cactus Subgraph[2]. All parameters keep being NP-complete to compute because the aforementioned base cases can be reduced to the general one by replacing each edge by one of multiplicity $r - 1$.

From the approximation point of view, it was proven in [FLMS12] that, when H is a planar graph, there is a randomized polynomial $O(1)$-approximation algorithm for $\mathsf{v\text{-}cover}_H$. For the cases of $\mathsf{v\text{-}cover}_{\mathcal{C}_{\theta_r}}$ and $\mathsf{v\text{-}pack}_{\mathcal{C}_{\theta_r}}$, $O(\log n)$-approximations are known for every $r \geq 1$ because of [JPS+14] (see also [SV05]). Moreover, $\mathsf{v\text{-}cover}_{\mathcal{C}_{\theta_r}}$ admits a deterministic 9-approximation [FJP10]. About $\mathsf{e\text{-}pack}_{\mathcal{C}_{\theta_r}}(G)$ it is known, from [KNY05], that there is a polynomial $O(\sqrt{\log n})$-approximation algorithm for the case where $r = 2$. Notice also that for $r = 1$, it is trivial to compute $\mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_r}}(G)$ in polynomial time. However, to our knowledge, nothing is known about the computation of $\mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_r}}(G)$ for $r \geq 3$.

Our results. In this paper we introduce a unified approach for the study of the combinatorial interconnections and the approximability of the parameters $\mathsf{v\text{-}cover}_{\mathcal{C}_{\theta_r}}$,

[2]The Maximum Cactus Subgraph problem asks, given a graph G and an integer k, whether G contains a subgraph with k edges where no two cycles share an edge. It is easy to reduce to this problem the Vertex Cycle Packing problem on cubic graphs which, in turn, can be proved to be NP-complete using a simple variant of the NP-completeness proof of Exact Cover by 2-Sets [Gol15].

$\mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_r}}$, $\mathsf{v\text{-}pack}_{\mathcal{C}_{\theta_r}}$, and $\mathsf{e\text{-}pack}_{\mathcal{C}_{\theta_r}}$. Our main combinatorial result is the following.

Theorem 1. *For every $r \in \mathbb{N}_{\geq 2}$ and every $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$ the graph class $\mathcal{C}_{\theta_r}$ has the* x-EP-property *with (optimal) gap function $f(k) = O(k \cdot \log k)$.*

Our proof is unified and treats simultaneously the covering and the packing parameters for both the vertex and the edge cases. This verifies the optimal combinatorial bound for the case where $\mathsf{x} = \mathsf{v}$ [FJS13] and optimally improves the previous bound in [RST13] for the case where $\mathsf{x} = \mathsf{e}$. Based on the proof of Theorem 1, we prove the following algorithmic result.

Theorem 2. *For every $r \in \mathbb{N}_{\geq 2}$ and every $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, there exists an $O(n \cdot \log(n) \cdot m)$-step algorithm that, given a graph G, outputs an $O(\log \mathrm{OPT})$-approximation for $\mathsf{x\text{-}cover}_{\mathcal{C}_{\theta_r}}$ and $\mathsf{x\text{-}pack}_{\mathcal{C}_{\theta_r}}$.*

Theorem 2 improves the results in [JPS+14] for the cases of $\mathsf{v\text{-}cover}_{\mathcal{C}_{\theta_r}}$ and $\mathsf{v\text{-}pack}_{\mathcal{C}_{\theta_r}}$ and, to our knowledge, this is the first approximation algorithm for $\mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_r}}$ and $\mathsf{e\text{-}pack}_{\mathcal{C}_{\theta_r}}$ for $r \geq 3$. We were also able to derive the following Erdős-Pósa-type result with linear gap on graphs of bounded tree-partition width (the definition of this width parameter is given in Subsection 2.2).

Theorem 3. *Let $t \in \mathbb{N}$. For every $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, the following holds: if $\mathcal{H}$ is a finite collection of connected graphs and G is a graph of tree-partition width at most t, then $\mathsf{x\text{-}cover}_{\mathcal{H}}(G) \leq \alpha \cdot \mathsf{x\text{-}pack}_{\mathcal{H}}(G)$, where α is a constant which depends only on t and $\mathcal{H}$.*

Let $\theta_{r,r'}$ (for some $r, r' \in \mathbb{N}_{\geq 1}$) denote the graph obtained by taking the disjoint union of θ_r and $\theta_{r'}$ and identifying one vertex of θ_r with one of $\theta_{r'}$. Another consequence of our results is that, for every $r, r' \in \mathbb{N}_{\geq 1}$, the class $\mathcal{C}_{\theta_{r,r'}}$ has the edge Erdős–Pósa property.

Theorem 4. *For every $r, r' \in \mathbb{N}$, there is a function $f_1^{r,r'} \colon \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$ such that for every simple graph G where $k = \mathsf{e\text{-}pack}_{\mathcal{C}_{\theta_{r,r'}}}(G)$, it holds that $\mathsf{e\text{-}cover}_{\mathcal{C}_{\theta_{r,r'}}}(G) \leq {f_1}^{r,r'}(k)$.*

Our techniques. Our proofs are based on the notion of *partitioned protrusion* that, roughly, is a tree-structured subgraph of G with small boundary to the rest of G (see Subsection 2.2 for the precise definition). Partitioned protrusions where essentially introduced in [CRST15] by the name *edge-protrusions* and can be seen as the edge-analogue of the notion of protrusions introduced in [BFL+09a] (see also [BFL+09b]). Our approach makes strong use of the main result of [CRST15], that is equivalently stated as Proposition 1 in this paper. According to this result, there exists a polynomial algorithm that, given a graph G and an integer k as an input, outputs one of the following: (1) a collection of k edge/vertex disjoint $\mathcal{C}_{\theta_r}$-subgraphs of G, (2) a $\mathcal{C}_{\theta_r}$-subgraph J with $O(\log k)$ edges, or (3) a large partitioned protrusion of G.

Our approximation algorithm does the following for each $k \leq |V(G)|$. If the first case of the above combinatorial result applies on G, we can safely output a packing of k

$\mathcal{C}_{\theta_r}$-subgraphs in G. In the second case, we make some progress as we may remove the vertices/edges of J from G and then set $k := k-1$. In order to deal with the third case, we prove that in a graph G with a sufficiently large partitioned protrusion, we can either find some $\mathcal{C}_{\theta_r}$-subgraph with $O(\log k)$ edges (which is the same as in the second case), or we can replace it by a smaller graph where both $\mathsf{x}\text{-}\mathsf{cover}_{\mathcal{H}}$ and $\mathsf{x}\text{-}\mathsf{pack}_{\mathcal{H}}$ remain invariant (Lemma 1). The proof that such a reduction is possible is given in Section 3 and is based on a suitable dynamic programming encoding of partial packings and coverings that is designed to work on partitioned protrusions.

Notice that the "essential step" in the above procedure is the second case that reduces the packing number of the current graph by 1 to the price of reducing the covering number by $O(\log k)$. This is the main argument that supports the claimed $O(\log \mathrm{OPT})$-approximation algorithm (Theorem 2) and the corresponding Erdős–Pósa relations in Theorem 1. Finally, Theorem 3 is a combinatorial implication of Lemma 1 and Theorem 4 follows by combining Theorem 3 with the results of Ding and Oporowski in [DO96].

Organization of the paper. In Section 2 we provide all concepts and notation that we use in our proofs. Section 3 contains the proof of Lemma 1, which is the main technical part of the paper. The presentation and analysis of our approximation algorithm is done in Section 4, where Theorem 1 and Theorem 2 are proven. Section 5 contains the proofs of Theorem 3 and Theorem 4. Finally, we summarize our results and provide several directions for further research in Section 6.

2 Preliminaries

2.1 Basic definitions

Let $\mathbf{t} = (x_1, \ldots, x_l) \in \mathbb{N}$ and $\chi, \psi : \mathbb{N} \to \mathbb{N}$. We say that $\chi(n) = O_{\mathbf{t}}(\psi(n))$ if there exists a computable function $\phi : \mathbb{N}^l \to \mathbb{N}$ such that $\chi(n) = O(\phi(\mathbf{t}) \cdot \psi(n))$.

Graphs. All graphs in this paper are undirected, loopless, and may have multiple edges. For this reason, a graph G is represented by a pair (V, E) where V is its vertex set, denoted by $V(G)$ and E is its edge multi-set, denoted by $E(G)$. The *edge-multiplicity* of an edge of G is the number of times it appears in $E(G)$. We set $n(G) = |V(G)|$ and $m(G) = |E(G)|$.If $\mathcal{H}$ is a finite collection of connected graphs, we set $n(\mathcal{H}) = \sum_{H\in\mathcal{H}} n(H)$, $m(\mathcal{H}) = \sum_{H\in\mathcal{H}} m(H)$, and $\bigcup\mathcal{H} = \bigcup_{G\in\mathcal{H}} G$.

Let $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$ where, in the rest of this paper, v will be interpreted as "vertex" and e will be interpreted as "edge". Given a graph G, we denote by $A_{\mathsf{x}}(G)$ the set of vertices or edges of G depending on whether $\mathsf{x} = \mathsf{v}$ or $\mathsf{x} = \mathsf{e}$, respectively.

Given a graph H and a graph J that are both subgraphs of the same graph G, we define the subgraph $H \cap_G J$ of G as the graph $(V(H) \cap V(J), E(H) \cap E(J))$.

Given a graph G and a set $S \subseteq V(G)$, such that all vertices in S have degree 2 in G, we define $\mathsf{diss}(G, S)$ as the graph obtained from G after we dissolve in it all vertices in S, i.e., replace each maximal path whose internal vertices are in S with an edge whose endpoints are the endpoints of the path.

Minors and topological minors. Given two graphs G and H, we say that H is a *minor* of G if there exits some function $\phi : V(H) \to 2^{V(G)}$ such that

- for every $v \in V(H)$, $G[\phi(v)]$ is connected;
- for every two distinct $v, u \in V(H)$, $\phi(v) \cap \phi(u) = \emptyset$; and
- for every edge $e = \{v, u\} \in E(H)$ of multiplicity l, there are at least l edges in G with one endpoint in $\phi(v)$ and the other in $\phi(u)$.

We say that H is a *topological minor* of G if there exits some collection $\mathcal{P}$ of paths in G and an injection $\phi : V(H) \to V(G)$ such that

- no path in $\mathcal{P}$ has an internal vertex that belongs to some other path in $\mathcal{P}$;
- $\phi(V(H))$ is the set of endpoints of the paths in $\mathcal{P}$; and
- for every two distinct $v, u \in V(H)$, $\{v, u\}$ is an edge of H of multiplicity l if and only if there are l paths in $\mathcal{P}$ between $\phi(v)$ and $\phi(u)$.

Given a graph H, we define by $\mathsf{ex}(H)$ the set of all topologically-minor minimal graphs that contain H as a minor. Notice that the size of $\mathsf{ex}(H)$ is upper-bounded by some function of $m(H)$ and that H is a minor of G if it contains a member of $\mathsf{ex}(H)$ as a topological minor. An *H-minor model of G* is every minimal subgraph of G that contains a member of $\mathsf{ex}(H)$ as a topological minor.

Packings and coverings. If G is a graph and $\mathcal{H}$ is a finite collection of connected graphs, an *$\mathcal{H}$-model* of G is a subgraph M of G that is a subdivision of a graph, denoted by $\hat{M}$, that is isomorphic to a member of $\mathcal{H}$. Clearly, the vertices of $\hat{M}$ are vertices of G and its edges correspond to paths in G between their endpoints such that internal vertices of a path do not appear in any other path. We refer to the vertices of $\hat{M}$ in G as the *branch vertices* of the $\mathcal{H}$-model M, whereas internal vertices of the paths between branch vertices will be called *subdivision vertices* of M. A graph which contains no $\mathcal{H}$-model is said to be $\mathcal{H}$-free. Notice that G has an $\mathcal{H}$-model iff G contains a graph of $\mathcal{H}$ as a topological minor.

An *x-$\mathcal{H}$-packing* of a graph G is a collection $\mathcal{P}$ of pairwise x-disjoint $\mathcal{H}$-models of G. Given an $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, we define $\mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(G)$ as the set of all x-$\mathcal{H}$-packings of G of size at least k. An *x-$\mathcal{H}$-covering* of a graph G is a set $C \subseteq A_{\mathsf{x}}(G)$ such that $G \setminus C$ does not contain

any $\mathcal{H}$-model. We define $\mathbf{C}^{\leq k}_{\mathsf{x},\mathcal{H}}(G)$ as the set of all x-H-coverings of G of size at most k. We finally define

$$\mathsf{x\text{-}cover}_{\mathcal{H}}(G) = \min\{k \mid \mathbf{C}^{\leq k}_{\mathsf{x},\mathcal{H}}(G) \neq \emptyset\}$$

and

$$\mathsf{x\text{-}pack}_{\mathcal{H}}(G) = \max\{k \mid \mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(G) \neq \emptyset\}.$$

It is easy to observe that, for every graph G and every finite collection of connected graphs $\mathcal{H}$, the following inequalities hold

$$\mathsf{v\text{-}cover}_{\mathcal{H}}(G) \leq \mathsf{e\text{-}cover}_{\mathcal{H}}(G), \qquad \mathsf{v\text{-}pack}_{\mathcal{H}}(G) \leq \mathsf{e\text{-}pack}_{\mathcal{H}}(G),$$
$$\mathsf{v\text{-}pack}_{\mathcal{H}}(G) \leq \mathsf{v\text{-}cover}_{\mathcal{H}}(G), \qquad \mathsf{e\text{-}pack}_{\mathcal{H}}(G) \leq \mathsf{e\text{-}cover}_{\mathcal{H}}(G).$$

2.2 Boundaried graphs

Informally, a boundaried graph will be used to represent a graph that has been obtained by "dissecting" a larger graph along some of its edges, where the boundary vertices correspond to edges that have been cut.

Boundaried graphs A *boundaried graph* $\mathbf{G} = (G, B, \lambda)$ is a triple consisting of a graph G, where B is a set of vertices of degree one (called *boundary*) and λ is a bijection from B to a subset of $\mathbb{N}_{\geq 1}$. The edges with at least one endpoint in B are called *boundary edges*. We define $E^{\mathsf{s}}(G)$ as the subset of $E(G)$ of boundary edges. We stress that instead of $\mathbb{N}_{\geq 1}$ we could choose any other set of symbols to label the vertices of B. We denote the set of labels of $\mathbf{G}$ by $\Lambda(\mathbf{G}) = \lambda(B)$. Given a finite collection of connected graphs, we say that a $\mathbf{G}$ is $\mathcal{H}$-free if $G \setminus B$ is $\mathcal{H}$-free.

Two boundaried graphs $\mathbf{G}_1$ and $\mathbf{G}_2$ are *compatible* if $\Lambda(\mathbf{G}_1) = \Lambda(\mathbf{G}_2)$. Let now $\mathbf{G}_1 = (G_1, B_1, \lambda_1)$ and $\mathbf{G}_2 = (G_2, B_2, \lambda_2)$ be two compatible boundaried graphs. We define the graph $\mathbf{G}_1 \oplus \mathbf{G}_2$ as the graph obtained by first taking the disjoint union of G_1 and G_2, then, for every $i \in \Lambda(\mathbf{G}_1)$, identifying $\lambda_1^{-1}(i)$ with $\lambda_2^{-1}(i)$, and finally dissolving all resulting identified vertices. Suppose that e is an edge of $G = \mathbf{G}_1 \oplus \mathbf{G}_2$ that was created after dissolving the vertex resulting from the identification of a vertex v_1 in B_1 and a vertex v_2 in B_2 and that e_i is the boundary edge of G_i that has v_i as endpoint, for $i = 1, 2$. Then we say that e is the *heir* of e_i in G, for $i = 1, 2$, and we denote this by $\mathsf{heir}_G(e_i)$. For $i \in \{1, 2\}$, if $S \subseteq E(G_i)$, then

$$\mathsf{heir}_G(S) = (E(G_i) \cap S) \cup \{\mathsf{heir}_G(e) \mid e \in E^{\mathsf{s}}(G_i) \cap S\}.$$

For reasons of notational consistency, if $V \subseteq V(G_i)$, we denote $\mathsf{heir}_G(S) = S$.

Figure 1 shows the result of the operation $\oplus$ on two graphs. Boundaries are drawn in gray and their labels are written next to them. The graphs $\mathbf{G}_1$ and $\mathbf{G}_2$ on this picture are compatible as $\Lambda(\mathbf{G}_1) = \Lambda(\mathbf{G}_2) = \{0, 1, 2, 3\}$.

For every $t \in \mathbb{N}_{\geq 1}$, we denote by $\mathcal{B}_t$ all boundaried graphs whose boundary is labeled by numbers in $[\![1, t]\!]$. Given a boundaried graph $\mathbf{G} = (G, B, \lambda)$ and a subset S of

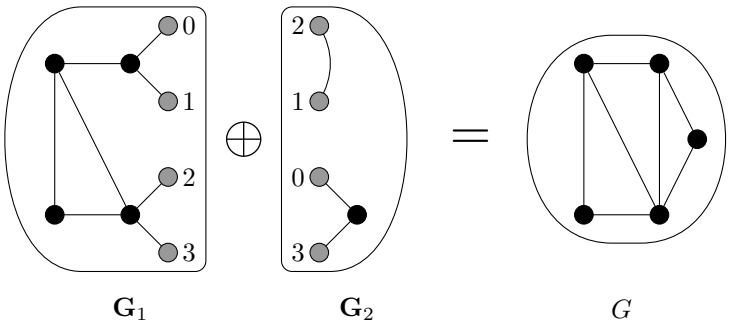

Figure 1: Gluing graphs together: $G = \mathbf{G}_1 \oplus \mathbf{G}_2$.

$V(G)$ such that all vertices in S have degree 2 in G, we define $\mathsf{diss}(\mathbf{G}, S)$ as the graph $\hat{\mathbf{G}} = (\hat{G}, B, \lambda)$ where $\hat{G} = \mathsf{diss}(G, S)$.

Let W be a graph and S be a non-empty subset of $V(W)$. An *S-splitting* of W is a pair $(\mathbf{G}_S, \mathbf{G}_{S^c})$ consisting of two boundaried graphs $\mathbf{G}_S = (G_S, B_S, \lambda_S)$ and $\mathbf{G}_{S^c} = (G_{S^s}, B_{S^s}, \lambda_{S^s})$ that can be obtained as follows: First, let W^+ be the graph obtained by subdividing in W every edge with one endpoint in S and the other in $V(W) \setminus S$ and let B be the set of created vertices. Let λ be any bijection from B to a subset of $\mathbb{N}_{\geq 1}$. Then $G_S = W^+[S \cup B]$, $G_{S^c} = W^+ \setminus S$, $B_S = B_{S^c} = B$, and $\lambda_S = \lambda_{S^c} = \lambda$. Notice that there are infinite such pairs, depending on the numbers that will be assigned to the boundaries of $\mathbf{G}_S$ and $\mathbf{G}_{S^c}$. Moreover, keep in mind that all the boundary edges of G_S are non-loop edges with exactly one endpoint in B and the same holds for the boundary edges of G_{S^s}. An example of a splitting is given in Figure 2, where boundaries are depicted by gray vertices.

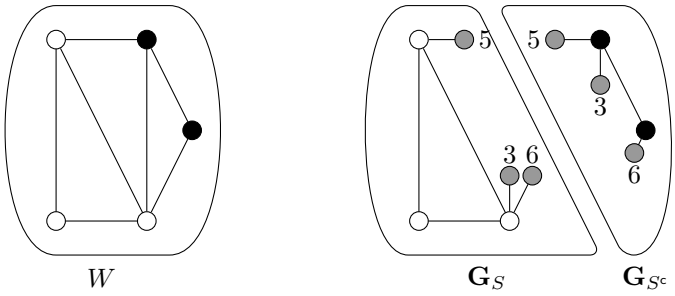

Figure 2: Cutting a graph: $(\mathbf{G}_S, \mathbf{G}_{S^c})$ is an S-splitting of W, where S consists of all the white vertices.

We say that $\mathbf{G}' = (G', B', \lambda')$ is a *boundaried subgraph* of $\mathbf{G} = (G, B, \lambda)$ if G' is a subgraph of G, $B' \subseteq B$ and $\lambda' = \lambda_{|B'}$. On the other hand, $\mathbf{G}$ is a subgraph of a (non-boundaried) graph H if $G = \mathbf{H}_S$ for some S-splitting $(\mathbf{H}_S, \mathbf{H}_{S^c})$, where $S \subseteq V(H)$.

If H is a graph, G is a subgraph of H, and $\mathbf{F} = (F, B, \lambda)$ is a boundaried subgraph

of H, we define $G \cap_H \mathbf{F}$ as follows. Let $S = V(G) \cap (V(F) \setminus B)$ and let G^+ be the graph obtained by subdividing once every edge of G that has one endpoint in S and the other in $V(G) \setminus S$. We call B' the set of created vertices and let $G' = G^+[S \cup B']$. Then G' is a subgraph of F where $B' \subseteq B$. For every $v \in B'$, we set $\lambda'(v) = \lambda(v)$, which is allowed according to the previous remark. Then $G \cap_H \mathbf{F} = (G', B', \lambda')$. Observe that $G \cap_H \mathbf{F}$ is an S-splitting of G.

Given two boundaried graphs $\mathbf{G}' = (G', B', \lambda')$ and $\mathbf{G} = (G, B, \lambda)$, we say that they are *isomorphic* if there is an isomorphism from G' to G that respects the labelings of B and B', i.e., maps every vertex $x \in B'$ to $\lambda^{-1}(\lambda'(x)) \in B$. Given a boundaried graph $\mathbf{G} = (G, B, \lambda)$, we denote $n(\mathbf{G}) = n(G) - |B|$ and $m(\mathbf{G}) = m(G)$.

Given a boundaried graph $\mathbf{G} = (G, B, \lambda)$ and an $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, we set $A_{\mathsf{x}}(\mathbf{G}) = V(G) \setminus B$ or $A_{\mathsf{x}}(\mathbf{G}) = E(G)$, depending on whether $\mathsf{x} = \mathsf{v}$ or $\mathsf{x} = \mathsf{e}$.

Partial structures. Let $\mathcal{F}$ be a family of graphs. A boundaried subgraph $\mathbf{J}$ of a boundaried graph $\mathbf{G}$ is a *partial $\mathcal{F}$-model* if there is a boundaried graph $\mathbf{H}$ which is compatible with $\mathbf{G}$ and a boundaried subgraph $\mathbf{J}'$ of $\mathbf{H}$ which is compatible with $\mathbf{J}$ such that $\mathbf{J} \oplus \mathbf{J}'$ is an $\mathcal{F}$-model of $\mathbf{G} \oplus \mathbf{H}$. Intuitively, this means that $\mathbf{J}$ can be extended into an $\mathcal{H}$-model in some larger graph. In this case, the $\mathcal{F}$-model $\mathbf{J} \oplus \mathbf{J}'$ is said to be an *extension* of $\mathbf{J}$.

Similarly, for every $p \in \mathbb{N}_{\geq 1}$, a collection of boundaried subgraphs $\mathcal{J} = \{\mathbf{J}_1, \ldots, \mathbf{J}_p\}$ of a graph $\mathbf{G}$ is a *partial* x-$\mathcal{F}$-*packing* if there is a boundaried graph $\mathbf{H}$ which is compatible with $\mathbf{G}$ and a collection of boundaried subgraphs $\{\mathbf{J}'_1, \ldots, \mathbf{J}'_p\}$ of $\mathbf{H}$ such that $\{\mathbf{J}_1 \oplus \mathbf{J}'_1, \ldots, \mathbf{J}_p \oplus \mathbf{J}'_p\}$ is an x-$\mathcal{F}$-packing of $\mathbf{G} \oplus \mathbf{H}$. The obtained packing is said to be an *extension* of the partial packing $\mathcal{J}$. A partial packing is $\mathcal{F}$-*free* if none of its members has an $\mathcal{F}$-model. Observe that every partial model of an $\mathcal{F}$-free partial packing in $\mathbf{G}$ must contain at least one boundary vertex of $\mathbf{G}$.

Partitions and protrusions. A *rooted tree-partition* of a graph G is a triple $\mathcal{D} = (\mathcal{X}, T, s)$ where (T, s) is a rooted tree and $\mathcal{X} = \{X_t\}_{t \in V(T)}$ is a partition of $V(G)$ where either $n(T) = 1$ or for every $\{x, y\} \in E(G)$, there exists an edge $\{t, t'\} \in E(T)$ such that $\{x, y\} \subseteq X_t \cup X_{t'}$ (see also [See85, Hal91, DO96]). Given an edge $f = \{t, t'\} \in E(T)$, we define E_f as the set of edges of G with one endpoint in X_t and the other in $X_{t'}$. Notice that all edges in E_f are non-loop edges. The *width* of $\mathcal{D}$ is defined as

$$\max\{|X_t|_{t \in V(T)}\} \cup \{|E_f|_{f \in E(T)}\}.$$

The *tree-partition width* of G is the minimum width over all tree-partitions of G and will be denoted by $\mathbf{tpw}(G)$. The elements of $\mathcal{X}$ are called *bags*.

In order to decompose graphs along edge cuts, we introduce the following edge-counterpart of the notion of *(vertex-)protrusion* introduced in [BFL+09a, BFL+09b].

Given a rooted tree-partition $\mathcal{D} = (\mathcal{X}, T, s)$ of G and a vertex $i \in V(T)$, we define

$$T_i = T[\mathsf{descendants}_{T,s}(i)], \quad V_i = \bigcup_{h \in V(T_i)} X_h, \quad \text{and} \quad G_i = G[V_i].$$

Let W be a graph and $t \in \mathbb{N}_{\geq 1}$. A pair $\mathbf{P} = (\mathbf{G}, \mathcal{D})$ is a *t-partitioned protrusion* of W if there exists an $S \subseteq V(W)$ such that

- $\mathbf{G} = (G, B, \lambda)$ is a boundaried graph where $\mathbf{G} \in \mathcal{B}_t$ and $\mathbf{G} = \mathbf{G}_S$ for some S-splitting $(\mathbf{G}_S, \mathbf{G}_{S^c})$ of W and
- $\mathcal{D} = (\mathcal{X}, T, s)$ is a rooted tree-partition of $G \setminus B$ of width at most t, where X_s are the neighbors in G of the vertices in B.

We say that a t-partitioned protrusion $(\mathbf{G}, \mathcal{D})$ of a graph W is *$\mathcal{H}$-free* if $\mathbf{G}$ is $\mathcal{H}$-free. For every vertex $u \in V(T)$, we also define the *t-partitioned protrusion $\mathbf{P}_u$ of W* as a pair $\mathbf{P}_u = (\mathbf{G}_u, \mathcal{D}_u)$, where $\mathcal{D}_u = (\{X_v\}_{v \in V_u}, T_u, u)$ and $\mathbf{G}_u = \mathbf{G}_{V_u}$ for some V_u-splitting $(\mathbf{G}_{V_u}, \mathbf{G}_{V_u^c})$ of W. We choose the labeling function of $\mathbf{G}_s$ so that it is the same as the one of $\mathbf{G}$, i.e., $\mathbf{G}_s = \mathbf{G}$. Notice that the labelings of all other $\mathbf{G}_u$'s are arbitrary. For every $u \in V(T)$ we define

$$\mathcal{G}_u = \{\mathbf{G}_l\}_{l \in \mathsf{children}_{(T,s)}(u)}.$$

2.3 Encodings, signatures, and folios

In this section we introduce tools that we will use to sort boundaried graphs depending on the models that are realizable inside.

Encodings. Let $\mathcal{H}$ be a family of graphs, let $t \in \mathbb{N}_{\geq 1}$, and let $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$. If $\mathbf{G} = (G, B, \lambda) \in \mathcal{B}_t$ is a boundaried graph and $S \subseteq A_{\mathsf{x}}(\mathbf{G})$, we define $\mu^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ as the collection of all sets $\{(\mathbf{J}_1, L_1), \ldots, (\mathbf{J}_\sigma, L_\sigma)\}$ such that

(i) $\{\mathbf{J}_1, \ldots, \mathbf{J}_\sigma\}$ is a partial x-$\mathcal{H}$-model of $G \setminus S$ of size σ and

(ii) $L_i = V(\hat{M}_i) \cap V(G)$, where M_i is an extension of J_i, for every $i \in [\![1, \sigma]\!]$.

In other words, L_i contains branch vertices of the partial model $\mathbf{J}_i$ for every $i \in [\![1, \sigma]\!]$ (see Figures 3 and 4). The set $\mu^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ encodes all different restrictions in $\mathbf{G}$ of partial x-$\mathcal{H}$-packings that avoid the set S. Given a boundaried graph $\mathbf{G} = (G, B, \lambda)$ and a set $L \subseteq V(G)$ such that every vertex of $V(G) \setminus L$ has degree 2 in G, we define $\kappa(\mathbf{G}, L)$ as the boundaried graph obtained from $\mathbf{G}$ by dissolving every vertex of $V(G) \setminus L$, i.e., $\kappa(\mathbf{G}, L) = (\mathsf{diss}(\mathbf{G}, V(G) \setminus L), B, \lambda)$. In the definition of κ we assume that the boundary vertices of $\kappa(\mathbf{G}, L)$ remain the same as in $\mathbf{G}$ while the other vertices are treated as new vertices (see Figure 5).

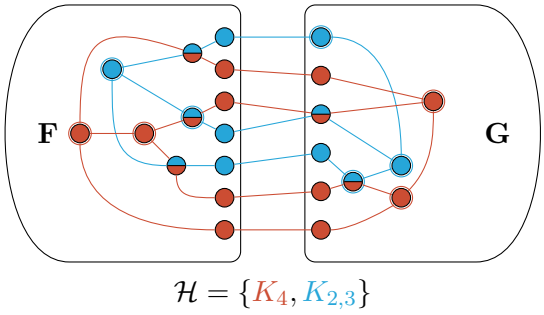

Figure 3: A member of $\mathbf{P}^{\geq 1}_{\mathsf{e},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G})$. Branch vertices are circled.

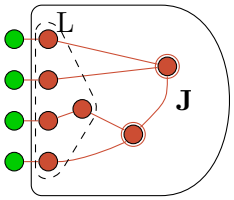

Figure 4: A partial model from the packing of Figure 3, where L is the set of subdivision vertices.

This allows us to introduce the following notation aimed at representing, intuitively, the essential part of each partial packing.

$$\begin{aligned}\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S) \quad = \quad & \{\hat{\mathcal{J}} = \{\hat{\mathbf{J}}_1, \ldots, \hat{\mathbf{J}}_\sigma\} = \{\kappa(\mathbf{J}_1, L_1), \ldots, \kappa(\mathbf{J}_\sigma, L_\sigma)\} \mid \\ & \{(\mathbf{J}_1, L_1), \ldots, (\mathbf{J}_\sigma, L_\sigma)\} \in \mu_{\mathcal{H}}(\mathbf{G}, S)\}.\end{aligned}$$

Isomorphisms. If $\mathbf{G} = (G, B, \lambda)$ and $\mathbf{G}' = (G', B', \lambda')$ are two compatible boundaried graphs in $\mathcal{B}_t$, $S \in V(G)$, and $S' \in V(G')$, we say that a member $\hat{\mathcal{J}}$ of $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ and a member $\hat{\mathcal{J}}'$ of $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', S')$ are *isomorphic* if there is a bijection between them such that paired elements are isomorphic. We also say that $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ and $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', S')$ are *isomorphic* if there is a bijection between them such that paired elements are isomorphic.

We now come to the point where we can define, for every boundaried graph, a signature encoding all the possible partial packings that can be realized in this graph.

Signatures and folios. For every $y \in \mathbb{N}$, we set

$$\mathsf{sig}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, y) \quad = \quad \{\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S),\ S \subseteq A_{\mathsf{x}}(G),\ |S| = y\}$$

and, given two compatible t-boundaried graphs $\mathbf{G}$ and $\mathbf{G}'$ and a $y \in \mathbb{N}$, we say that $\mathsf{sig}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, y)$ and $\mathsf{sig}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', y)$ are *isomorphic* if there is a bijection between them such that paired elements are isomorphic.

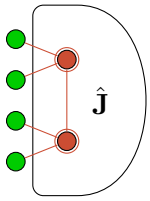

Figure 5: The compression of the partial packing of Figure 4: $\hat{\mathbf{J}} = \kappa(\mathbf{J}, L)$.

Finally, for $\rho \in \mathbb{N}$, we set

$$\mathsf{folio}_{\mathcal{H},\rho}(\mathbf{G}) \quad = \quad (\mathsf{sig}^{\mathsf{v}}_{\mathcal{H}}(\mathbf{G},0),\ldots,\mathsf{sig}^{\mathsf{v}}_{\mathcal{H}}(\mathbf{G},\rho),\mathsf{sig}^{\mathsf{e}}_{\mathcal{H}}(\mathbf{G},0),\ldots,\mathsf{sig}^{\mathsf{e}}_{\mathcal{H}}(\mathbf{G},\rho)).$$

Given two t-boundaried graphs $\mathbf{G}$ and $\mathbf{G}'$, a $\rho \in \mathbb{N}$, and a finite collection of connected graphs $\mathcal{H}$, we say that $\mathbf{G} \simeq_{\mathcal{H},\rho} \mathbf{G}'$ if $\mathbf{G}$ and $\mathbf{G}'$ are compatible, neither $\mathbf{G}$ nor $\mathbf{G}'$ contains an $\mathcal{H}$-model, and the elements of $\mathsf{folio}_{\mathcal{H},\rho}(\mathbf{G})$ and $\mathsf{folio}_{\mathcal{H},\rho}(\mathbf{G}')$ are position-wise isomorphic.

3 The reduction

The purpose of this section is to prove the following lemma.

Lemma 1. *There exists a function $f_2\colon \mathbb{N}^2 \to \mathbb{N}$ and an algorithm that, given a positive integer t, a finite collection $\mathcal{H}$ of connected graphs where $h = m(\mathcal{H})$, and a t-partitioned protrusion $\mathbf{P} = (\mathbf{G}, (\mathcal{X}, T, s))$ of a graph W with $n(\mathbf{G}) > f_2(h,t)$, outputs either*

- *an $\mathcal{H}$-model of W with at most $f_2(h,t)$ edges or*
- *a graph W' such that*

$$\begin{aligned} \mathsf{x\text{-}pack}_{\mathcal{H}}(W') &= \mathsf{x\text{-}pack}_{\mathcal{H}}(W), \\ \mathsf{x\text{-}cover}_{\mathcal{H}}(W') &= \mathsf{x\text{-}cover}_{\mathcal{H}}(W), \text{ and} \\ n(W') &< n(W). \end{aligned}$$

Furthermore, this algorithm runs in time $O_{t,h}(n(T))$.

In other words we can, in linear time, either find a *small* $\mathcal{H}$-model, or reduce the graph to a smaller one where the parameters of packing and covering stay the same. Before giving the proof of Lemma 1, we need to prove several intermediate results. In the sequel, unless stated otherwise, we assume that $\mathsf{x} \in \{\mathsf{v},\mathsf{e}\}$, $t \in \mathbb{N}_{\geq 1}$ and that $\mathcal{H}$ is a finite collection of connected graphs. We set $h = m(\mathcal{H})$.

Lemma 2. *There are two functions $f_3\colon \mathbb{N}^2 \to \mathbb{N}$ and $f_4\colon \mathbb{N}^2 \to \mathbb{N}$ such that, for every graph W and every t-partitionned protrusion $(\mathbf{G}, (\mathcal{X}, T, s))$ of W, if $\mathcal{P}$ is an $\mathcal{H}$-free partial x-$\mathcal{H}$-packing in $\mathbf{G}$ then:*

(a) The partial models of graphs in $\mathcal{H}$ that are contained in $\mathcal{P}$ have in total at most $f_3(h,t)$ branch vertices.

(b) $\mathcal{P}$ intersects at most $f_4(h,t)$ graphs of $\mathcal{G}_s$.

Proof. Proof of (a). First, note that any $\mathcal{H}$-free partial x-$\mathcal{H}$-packing in $\mathbf{G}$ has cardinality at most t, because each partial model it contains must use a boundary edge of $\mathbf{G}$, and two distinct models of the same packing are (at least) edge-disjoint. Also, each of these partial models contains at most $\max_{H\in\mathcal{H}} n(H) \leq h$ branch vertices. Consequently, for every $\mathcal{H}$-free partial x-$\mathcal{H}$-packing in $\mathbf{G}$, the number of branch vertices of graphs of $\mathcal{H}$ it induces in $\mathbf{G}$ is at most $t \cdot h$. Hence the function $f_3(h,t) := t \cdot h$ upper-bounds the amount of branch vertices each $\mathcal{H}$-free partial packing can contain.

Proof of (b). Let ζ be the maximum multiplicity of an edge in a graph of $\mathcal{H}$. Because of (a), every $\mathcal{H}$-free partial x-$\mathcal{H}$-packing $\mathcal{P}$ in $\mathbf{G}$ has at most $f_3(h,t)$ branch vertices of graphs of $\mathcal{H}$, so at most $f_3(h,t)$ graphs of $\mathcal{G}_s$ may contain such vertices. Besides, $\mathcal{P}$ might also contain paths free of branch vertices linking pairs of branch vertices. Since there are at most $(f_3(h,t))^2$ such pairs and no pair will need to be connected with more than $\zeta \leq h$ distinct paths, it follows that at most $(f_3(h,t))^2 \cdot h$ graphs of $\mathcal{G}_s$ contain vertices from these paths. Therefore, every $\mathcal{H}$-free partial x-$\mathcal{H}$-packing intersects at most $f_3(h,t) + (f_3(h,t))^2 \cdot h =: f_4(h,t)$ graphs of $\mathcal{G}_s$. □

Lemma 3. *There is a function $f_5\colon \mathbb{N}^2 \to \mathbb{N}$ such that the image of the function $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}$, when its domain is restricted to*

$$\{(\mathbf{G}, S),\ \mathbf{G} \text{ is } \mathcal{H}\text{-free and } S \subseteq A_{\mathsf{x}}(\mathbf{G})\},$$

has size upper-bounded by $f_5(h,t)$.

Proof. Let $\mathbf{G}$ be an $\mathcal{H}$-free t-boundaried graph and let $S \subseteq A_{\mathsf{x}}(\mathcal{G})$. By Lemma 2(a), every $\mathcal{H}$-free partial x-$\mathcal{H}$-packing in $\mathbf{G}$ contains at most $f_3(h,t)$ branch vertices. This partial packing can in addition use at most t boundary vertices. Let $\mathcal{C}_{h,t}$ be the class of all $(\leq t)$-boundaried graphs on at most $f_3(h,t) + t$ vertices. Clearly the size of this class is a function depending on h and t only. Recall that the elements of the set $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ are obtained from partial x-$\mathcal{H}$-packings by dissolving internal vertices of the paths linking branch vertices, hence every element of $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ is a t-boundaried graph of $\mathcal{B}_t$ having at most $f_3(h,t) + t$ vertices. Therefore, for any $\mathcal{H}$-free t-boundaried graph $\mathbf{G}$ and subset $S \subseteq A_{\mathsf{x}}(\mathbf{G})$, we have $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S) \subseteq \mathcal{C}_{h,t}$. As a consequence, the image of the function $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}$ when restricted to $\mathcal{H}$-free t-boundaried graphs $\mathbf{G} \in \mathcal{B}_t$ (and subsets $S \subseteq A_{\mathsf{x}}(\mathbf{G})$) is a subset of the power set of $\mathcal{C}_{h,t}$, so its size is upper-bounded by a function (which we call f_5) that depends only on h and t. □

Corollary 1. *There is a function $f_6\colon \mathbb{N}^2 \to \mathbb{N}$ such that the relation $\simeq_{\mathcal{H},t}$ partitions $\mathcal{H}$-free t-boundaried graphs into at most $f_6(h,t)$ equivalence classes.*

The following lemma follows directly from the definition of $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}$.

Lemma 4. *Let $\mathbf{F}, \mathbf{G} \in \mathcal{B}_t$ be two compatible t-boundaried graphs and let $k \in \mathbb{N}$. Then we have:*

$$\mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G}) \neq \emptyset \iff \exists \hat{\mathcal{J}} \in \hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, \emptyset),\ \mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{U}\hat{\mathcal{J}}) \neq \emptyset.$$

The choice of the definition of the relation $\simeq$ is justified by the following lemma. Roughly speaking, it states that we can replace a t-partitioned protrusion of a graph with any other $\simeq_{\mathcal{H},t}$-equivalent t-partitioned protrusion without changing the covering and packing number of the graph. The reduction algorithm that we give after this lemma relies on this powerful property.

Lemma 5 (protrusion replacement)**.** *Let $\mathbf{F}, \mathbf{G}, \mathbf{G}' \in \mathcal{B}_t$ be three compatible boundaried graphs such that $\mathbf{G} \simeq_{\mathcal{H},t} \mathbf{G}'$. For every $k \in \mathbb{N}$, we have:*

(i) $\mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G}) \neq \emptyset \iff \mathbf{P}^{\geq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G}') \neq \emptyset$ *and*

(ii) $\mathbf{C}^{\leq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G}) \neq \emptyset \iff \mathbf{C}^{\leq k}_{\mathsf{x},\mathcal{H}}(\mathbf{F} \oplus \mathbf{G}') \neq \emptyset$.

Proof. Proof of item (i), "⇒". Let $\mathcal{M}$ be an x-$\mathcal{H}$-packing of size at least k in $\mathbf{F} \oplus \mathbf{G}$, whose set of branch vertices is L. We define

$$\begin{aligned}
\mathbf{J}_{\mathbf{F}} &= (\mathbf{U}\mathcal{M}) \cap_{\mathbf{F}\oplus\mathbf{G}} \mathbf{F}, \\
\mathbf{J}_{\mathbf{G}} &= (\mathbf{U}\mathcal{M}) \cap_{\mathbf{F}\oplus\mathbf{G}} \mathbf{G}, \text{and} \\
\hat{\mathbf{J}}_{\mathbf{G}} &= \bigcup_{M \in \mathcal{M}} \kappa(M \cap_{\mathbf{F}\oplus\mathbf{G}} \mathbf{G}, L \cap V(G)).
\end{aligned}$$

Note that $\hat{\mathbf{J}}_{\mathbf{G}} \in \hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, \emptyset)$ and that $\mathbf{F} \oplus \hat{\mathbf{J}}_{\mathbf{G}}$ has an x-$\mathcal{H}$-packing of size at least k (cf. Lemma 4). By definition of $\simeq$, there is a bijection ψ between $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, \emptyset)$ and $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', \emptyset)$. Let $\hat{\mathbf{J}}'_{\mathbf{G}}$ be the image of $\hat{\mathbf{J}}_{\mathbf{G}}$ by ψ. Since $\hat{\mathbf{J}}'_{\mathbf{G}}$ and $\hat{\mathbf{J}}_{\mathbf{G}}$ are isomorphic, $\mathbf{F} \oplus \hat{\mathbf{J}}'_{\mathbf{G}}$ also has an x-$\mathcal{H}$-packing of size at least k. By Lemma 4, this implies that such a packing exists in $\mathbf{F} \oplus \mathbf{G}'$ as well. The direction "⇐" is symmetric as $\mathbf{G}$ and $\mathbf{G}'$ play the same role.

Proof of item (ii), "⇒". Let $C \subseteq A_{\mathsf{x}}(\mathbf{F} \oplus \mathbf{G})$ be a minimum x-$\mathcal{H}$-covering of $\mathbf{F} \oplus \mathbf{G}$ of size at most k. Let $S = C \cap A_{\mathsf{x}}(\mathbf{G})$. Since we assume that $\mathbf{G}$ is $\mathcal{H}$-free and that C is minimum, we can also assume that $|S| \leq t$ (otherwise we could get a smaller covering by taking the t boundary vertices/edges of $\mathbf{G}$). By our assumption that $\mathbf{G} \simeq_{\mathcal{H},t} \mathbf{G}'$, there is an isomorphism between $\mathsf{sig}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, |S|)$ and $\mathsf{sig}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', |S|)$. Let $S' \subseteq A_{\mathsf{x}}(\mathbf{G}')$ be a set such that $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ is sent to $\hat{\mu}^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}', S')$ by this isomorphism. Then observe that every partial packing $\mathcal{J}'$ of $\mathbf{G}' \setminus S'$, such that $(\mathbf{F} \setminus C) \oplus (\mathbf{U}\mathcal{J}')$ has an $\mathcal{H}$-model, can be translated into a partial packing $\mathcal{J}$ of $\mathbf{G} \setminus S$ such that $(\mathbf{F} \setminus C) \oplus (\mathbf{U}\mathcal{J})$ also has such a model, in the same way as in the proof of item (i) above. As C is a cover, this would lead to contradiction. Therefore $\mu^{\mathsf{x}}_{\mathcal{H}}(\mathbf{G}, S)$ does not contain such a partial packing. As a consequence, $C \cap A_{\mathsf{x}}(\mathbf{F}) \cup S'$ is a covering of $\mathbf{F} \oplus \mathbf{G}'$ of size at most k. As in the previous case, the proof of direction "⇐" comes from the symmetry in the statement. □

Lemma 5 can be rewritten as follows.

Corollary 2. *Under the assumptions of Lemma 5, we have* $\mathsf{x\text{-}pack}(\mathbf{F} \oplus \mathbf{G}) = \mathsf{x\text{-}pack}(\mathbf{F} \oplus \mathbf{G}')$ *and* $\mathsf{x\text{-}cover}(\mathbf{F} \oplus \mathbf{G}) = \mathsf{x\text{-}cover}(\mathbf{F} \oplus \mathbf{G}')$.

Let $f_7(h,t) = f_6(h,t) \cdot f_4(h,t)$ and $f_8(h,t) = (f_7(h,t))^{f_6(h,t)}$.

Lemma 6. *Let* $\mathbf{P} = (\mathbf{G}, (\mathcal{X}, T, s))$ *be a* t*-partitioned protrusion of a graph* W*, and let* $u \in V(T)$ *be a vertex with more than* $f_7(h,t)$ *children such that for every* $v \in \mathsf{children}_{(T,s)}(u)$*, we have* $m(\mathbf{G}_v) \leq f_8(h,t)$*. Then, either*

- W *contains an* $\mathcal{H}$*-model* M *with at most* $f_8(h,t)$ *edges or*
- *there exists a graph* W' *such that*

$$\begin{aligned} \mathsf{x\text{-}pack}_{\mathcal{H}}(W') &= \mathsf{x\text{-}pack}_{\mathcal{H}}(W), \\ \mathsf{x\text{-}cover}_{\mathcal{H}}(W') &= \mathsf{x\text{-}cover}_{\mathcal{H}}(W), \text{ and} \\ n(W') &< n(W). \end{aligned}$$

Moreover, there is an algorithm that, given such $\mathbf{P}, W$*, and* u*, returns either* M *or* W' *as above in* $O_{h,t}(1)$ *steps.*

Proof. As u has more than $f_7(h,t)$ children, it contains a collection of $d = f_4(h,t) + 1$ children $v_1, \ldots, v_d$, such that $\mathbf{G}_{v_1} \simeq_{\mathcal{H},t} \mathbf{G}_{v_i}$ for every $i \in [\![2, d]\!]$ (by the pigeonhole principle and since $\simeq_{\mathcal{H},t}$ has at most $f_6(h,t)$ equivalence classes). Let us now assume that $\mathbf{G}_u$ is $\mathcal{H}$-free. Since every x-$\mathcal{H}$-packing of W will touch at most $f_4(h,t)$ bags of children of u (by Lemma 2(b)), we can safely delete one of the $f_4(h,t) + 1$ equivalent subgraphs mentioned above. We use the following algorithm in order to find such a bag to delete or a small $\mathcal{H}$-model.

1. Let A be an array of $f_6(h,t)$ counters initialized to 0, each corresponding to a distinct equivalence class of $\simeq_{\mathcal{H},t}$;
2. Pick a vertex $v \in \mathsf{children}_{(T,s)}(u)$ that has not been considered yet;
3. If G_v contains an $\mathcal{H}$-model $\mathcal{M}$, then return $\mathcal{M}$ and exit;
4. Otherwise, increment the counter of A corresponding to the equivalence class of $\mathbf{G}_u$ by one;
5. If this counter reaches $d + 1$, return v, otherwise go back to Line 2.

Notice that the model returned in Line 3 has size at most $t \cdot f_8(h,t)$ (as we assume that $m(G_v) \leq f_8(h,t)$) and that the vertex returned in Line 5 has the desired property. By Corollary 1, the relation $\simeq_{\mathcal{H},t}$ has at most $f_6(h,t)$ equivalence classes, thus the main loop will be run at most $f_6(h,t) \cdot f_4(h,t) + 1$ times (by the pigeonhole principle). Eventually,

Lines 3 and 4 can be performed in $O_{h,t}(1)$-time given that G_v has size bounded by a function of h and t.

In the end, we return $W' = W \setminus V(G_v)$ if the algorithm outputs v and $\mathcal{M}$ otherwise. □

Lemma 7. *There is an algorithm that, given a t-partitioned protrusion $\mathbf{P} = (\mathbf{G}, (\mathcal{X}, T, s))$ of a graph W and a vertex $u \in V(T)$ such that*

- *u has height exactly $f_6(h,t)$ in (T,s),*
- *the graph of $\mathbf{G}_u$ is $\mathcal{H}$-free, and*
- *T_u has maximum degree at most $f_7(h,t)$,*

outputs a graph W' such that

$$\begin{aligned} \mathsf{x\text{-}pack}_{\mathcal{H}}(W') &= \mathsf{x\text{-}pack}_{\mathcal{H}}(W), \\ \mathsf{x\text{-}cover}_{\mathcal{H}}(W') &= \mathsf{x\text{-}cover}_{\mathcal{H}}(W), \text{ and} \\ n(W') &< n(W). \end{aligned}$$

Moreover, this algorithm runs in $O_{h,t}(1)$-time.

Proof. By definition of vertex u, there is a path of $f_6(h,t)+1$ vertices from a leaf of T_u to u. Let us arbitrarily choose, for every vertex v of this path, a V_v-splitting $(\mathbf{G}_{G_v}, \mathbf{G}_{G_v^c})$ of $\mathbf{G}$. According to Corollary 1, there are two distinct vertices v, w on this path such that $\mathbf{G}_v \simeq_{\mathcal{H},t} \mathbf{G}_w$. Since T_u has order bounded by a function of h and t, finding these two vertices can be done in $O_{h,t}(1)$-time. Let us assume without loss of generality that s is closer to v than w. Let $\mathbf{H}$ be the boundaried graph such that $W = \mathbf{H} \oplus \mathbf{G}_{G_v}$ and let $W' = \mathbf{H} \oplus \mathbf{G}_{G_w}$. By Corollary 2, we have $\mathsf{x\text{-}pack}_{\mathcal{H}}(W') = \mathsf{x\text{-}pack}_{\mathcal{H}}(W)$ and $\mathsf{x\text{-}cover}_{\mathcal{H}}(W') = \mathsf{x\text{-}cover}_{\mathcal{H}}(W)$. Furthermore, the graph W' is clearly smaller than W. □

For every $h, t \in \mathbb{N}^2$, let $f_2(h,t) = t \cdot f_8(h,t)$. We are now ready to prove Lemma 1.

Proof of Lemma 1. Observe that since $n(G) > f_2(h,t)$ and each bag of $(\mathcal{X}, T, s)$ contains at most t vertices, we have

$$n(T) > f_2(h,t)/t = (f_6(h,t) \cdot f_4(h,t))^{f_6(h,t)}.$$

Therefore, T has either diameter more than $f_6(h,t)$ or a vertex of degree more than $f_6(h,t) \cdot f_4(h,t)$.

Let us consider the following procedure.

1. By a BFS on (T,s), compute the height of each vertex of T and find (if it exists) a vertex v of degree more than $f_7(h,t)$ and height at most $f_6(h,t)-1$ that has minimum height.

2. If such a vertex v is found, then apply the algorithm of Lemma 6 on $\mathbf{P}$ and v, and return the obtained result.

3. Otherwise, find a vertex u of height exactly $f_6(h,t)$ in (T,s) and then apply the algorithm of Lemma 7 on $\mathbf{P}$ and (T_u,u) and return the obtained result.

The correctness of this algorithm follows from Lemma 6 and Lemma 7. The BFS done in the first step takes time $O(n(T))$ and the rest of the algorithm takes time $O_{h,t}(1)$ according to the aforementioned lemmata. □

4 From the Erdős-Pósa property to approximation

For the purposes of this section we define $\Theta_r = \mathsf{ex}(\theta_r)$. We need the following that is one of the main results in [CRST15, Theorem 3].

Proposition 1. *There is an algorithm that, with input three positive integers r, w, z and a graph W, outputs one of the following:*

- *a Θ_r-model of W with at most z edges,*
- *a $(2r-2)$-partitioned protrusion $(\mathbf{G},\mathcal{D})$ of W, where $\mathbf{G}=(G,B,\lambda)$ and such that G is a connected graph and $n(\mathbf{G})>w$, or*
- *an H-minor model of W for some graph H with $\delta(H)\geq \frac{1}{r-1}2^{\frac{z-5r}{4r(2w+1)}}$,*

in $O_r(m)$ steps.

4.1 A lemma on reduce or progress

The proof of the next lemma combines Proposition 1 and Lemma 1.

Lemma 8 (Reduce or progress)**.** *There is an algorithm that, with input $\mathsf{x}\in\{\mathsf{v},\mathsf{e}\}$, $r\in\mathbb{N}_{\geq 2}$, $k\in\mathbb{N}$ and an n-vertex graph W, outputs one of the following:*

- *a Θ_r-model of W with at most $O_r(\log k)$ edges;*
- *a graph W' where*

$$\begin{aligned}\mathsf{x\text{-}cover}_{\mathcal{H}}(W') &= \mathsf{x\text{-}cover}_{\mathcal{H}}(W),\\ \mathsf{x\text{-}pack}_{\mathcal{H}}(W') &= \mathsf{x\text{-}pack}_{\mathcal{H}}(W),\ \textit{and}\\ n(W') &< n(W);\textit{or}\end{aligned}$$

- *an H-minor model in W, for some graph H with $\delta(H)\geq k(r+1)$,*

in $O_r(m)$ steps.

Proof. We set $t = 2r-2$, $w = f_2(h,t)$, $z = 2r(w-1)\log(k(r+1)(r-1)) + 5r$, and $h = m(\Theta_r)$. Observe that $z = O_r(\log k)$ and $h, t, w = O_r(1)$. Also observe that our choice for variable z ensures that $2^{\frac{z-5r}{2r(w-1)}}/(r-1) = k(r+1)$.

By applying the algorithm of Proposition 1 to r, w, z, and W, we obtain in $O_r(m(W))$-time either:

First case: a Θ_r-model in W of at most z edges,

Second case: a $(2r-2)$-edge-protrusion Y of W with extension $> w$, or

Third case: an H-minor model M in W, for some graph H with $\delta(H) \geq k(r+1)$.

In the first case, we return the obtained Θ_r-model.

In the second case, by applying the algorithm of Lemma 1 on Y, we get in $O(n(W))$-time either a Θ_r-model of W on at most $w = O_r(1)$ vertices, or a graph W' where, for $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, $\mathsf{x}\text{-}\mathsf{cover}_{\mathcal{H}}(W') = \mathsf{x}\text{-}\mathsf{cover}_{\mathcal{H}}(W)$, $\mathsf{x}\text{-}\mathsf{pack}_{\mathcal{H}}(W') = \mathsf{x}\text{-}\mathsf{pack}_{\mathcal{H}}(W)$ and $n(W') < n(W)$.

In the third case, we return the model M.

In each of the above cases, we get after $O(m)$ steps either a model of a graph with minimum degree more than $k(r+1)$, a Θ_r-model in W with at most z edges, or an equivalent graph of smaller size. □

It might not be clear yet to what purpose the model of a graph of degree more than $k(r+1)$ output by the algorithm of Lemma 8 can be used. An answer is given by the following lemmata, which state that such a graph contains a packing of at least k models of Θ_r. These lemmata will be used in the design of the appoximation algorithms in Subsection 4.2.

Lemma 9. *There is an algorithm that, given $k, r \in \mathbb{N}_{\geq 1}$ and a graph G with $\delta(G) \geq kr$, returns a member of $\mathbf{P}^{\geq k}_{\mathsf{e},\Theta_r}(G)$ in G in $O(m)$ steps.*

Proof. Starting from any vertex u, we grow a maximal path P in G by iteratively adding to P a vertex that is adjacent to the previously added vertex but does not belong to P. Since $\delta(G) \geq kr$, any such path will have length at least $kr+1$. At the end, all the neighbors of the last vertex v of P belong to P (otherwise P could be extended). Since v has degree at least kr, v has at least kr neighbors in P. Let $w_0, \ldots, w_{kr-1}$ be an enumeration of the kr first neighbors of v in the order given by P, starting from u. For every $i \in [\![0, k-1]\!]$, let S_i be the subgraph of G induced by v and the subpath of P starting at w_{ir} and ending at $w_{(i+1)r-1}$. Observe that for every $i \in [\![0, k-1]\!]$, S_i contains a Θ_r-model and that the intersection of every pair of graphs from $\{S_i\}_{i \in [\![0,k-1]\!]}$ is $\{v\}$. Hence P contains a member of $\mathbf{P}^{\geq k}_{\mathsf{e},\Theta_r}(G)$, as desired. Every edge of G is considered at most once in this algorithm, yielding to a running time of $O(m)$ steps. □

Corollary 3. *There is an algorithm that, given $r \in \mathbb{N}_{\geq 1}$ and a graph G with $\delta(G) \geq r$, returns a Θ_r-model in G in $O(m)$-steps.*

Observe that the previous lemma only deals with edge-disjoint packings. An analogue of Lemma 9 for vertex-disjoint packings can be proved using Proposition 2, to the price of a worse time complexity.

Proposition 2 (Theorem 12 of [BTV07])**.** *Given $k, r \in \mathbb{N}_{\geq 1}$ and an input graph G such that $\delta(G) \geq k(r+1)-1$, a partition $(V_1, \ldots, V_k)$ of $V(G)$ satisfying $\forall i \in [\![1, k]\!]$, $\delta(G[V_i]) \geq r$ can be found in $O(n^c)$ steps for some $c \in \mathbb{N}$.*

Lemma 10. *There is an algorithm that, given $k, r \in \mathbb{N}_{\geq 1}$ and a graph G with $\delta(G) \geq k(r+1)-1$, outputs a member of $\mathbf{P}^{\geq k}_{\mathsf{v},\Theta_r}(G)$ in $O(n^c + m)$ steps, where c is the constant of Proposition 2.*

Proof. After applying the algorithm of Proposition 2 on G to obtain in $O(n^c)$-time k graphs $G[V_1], \ldots, G[V_k]$, we extract a Θ_r-model from each of them using Corollary 3. □

4.2 Approximation algorithms

Theorem 1 is a direct combinatorial consequence of the following.

Theorem 5. *There is a function $f_9 \colon \mathbb{N} \to \mathbb{N}$ and an algorithm that, with input $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, $r \in \mathbb{N}_{\geq 2}$, $k \in \mathbb{N}$, and an n-vertex graph W, outputs either a x-Θ_r-packing of W of size k or an x-Θ_r-covering of W of size at most $f_9(r) \cdot k \cdot \log k$. Moreover, this algorithm runs in $O(n \cdot m)$ steps if $\mathsf{x} = \mathsf{e}$ and in $O(n^c + n \cdot m)$ steps if $\mathsf{x} = \mathsf{v}$, where c is the constant from Proposition 2.*

Proof. Let $f : \mathbb{N} \to \mathbb{N}$ be a function such that each Θ_r-model output by the algorithm of Lemma 8 has size at most $f_9(r) \cdot \log k$. We consider the following procedure.

1. $G := W$; $P := \emptyset$;
2. Apply the algorithm of Lemma 8 on (x, r, k, G):

   Progress: if the output is a Θ_r-model M, let $G := G \setminus A_{\mathsf{x}}(M)$ and $P = P \cup \{M\}$;

   Win: if the output is a H-minor model M in W for some graph H with $\delta(H) \geq k(r+1)$, apply the algorithm of Lemma 9 (if $\mathsf{x} = \mathsf{e}$) or the one of Lemma 10 (if $\mathsf{x} = \mathsf{v}$) to H to obtain a member of $\mathbf{P}^{\geq k}_{\mathsf{x},\Theta_r}(H)$. Using M, translate this packing into a member of $\mathbf{P}^{\geq k}_{\mathsf{x},\Theta_r}(W)$ and return this new packing;

   Reduce: otherwise, the output is a graph G'. Then let $G := G'$.
3. If $|P| = k$ then return P which is a member of $\mathbf{P}^{\geq k}_{\mathsf{x},\Theta_r}(W)$;
4. If $n(W) = 0$ then return P which is in this case a member of $\mathbf{C}^{\leq f_9(r)\log k}_{\mathsf{x},\Theta_r}(W)$;

5. Otherwise, go back to Line 2.

This algorithm clearly returns the desired result. Furthermore, the loop is executed at most $n(W)$ times and each call to the algorithm of Lemma 8 takes $O(m(W))$ steps. When the algorithm reaches the "Win" case (which can happen at most once), the calls to the algorithm of Lemma 9 (if $\mathsf{x} = \mathsf{e}$) or the one of Lemma 10 (if $\mathsf{x} = \mathsf{v}$), respectively, take $O(m(H))$ and $O\left((n(H))^c\right)$ steps. Therefore, in total, this algorithm terminates in $O(n \cdot m)$ steps if $\mathsf{x} = \mathsf{e}$ and in $O\left(n^c + n \cdot m\right)$ steps if $\mathsf{x} = \mathsf{v}$. □

Observe that if the algorithm of Theorem 5 reaches the "Win" case, then the input graph is known to contain an x-Θ_r-packing of size at least k. As a consequence, if we are only interested in the existence of a packing or covering, the call to the algorithm of Lemma 9 or Lemma 10 is not necessary.

Corollary 4. *There is an algorithm that, with input $\mathsf{x} \in \{\mathsf{v}, \mathsf{e}\}$, $r \in \mathbb{N}_{\geq 2}$, $k \in \mathbb{N}$, and a graph W, outputs 0 only if W has an x-Θ_r-packing of size k or 1 only if W has an x-Θ_r-covering of size at most $f_9(r) \cdot k \cdot \log k$. Furthermore this algorithm runs in $O(n \cdot m)$ steps.*

We now conclude this section with the proof of Theorem 2.

Proof of Theorem 2. Let us call A the algorithm of Corollary 4. Let $k_0 \in [\![1, n(W)]\!]$ be an integer such that $\mathsf{A}((\mathsf{x}, r, k_0, W) = 1$ and $\mathsf{A}(\mathsf{x}, r, k_0 - 1, W) = 0$, and let us show that the value $k_0 \log k_0$ is an $O(\log OPT)$-approximation of $\mathbf{p}(W)$.

First, notice that for every $k > \mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W)$, the value returned by $\mathsf{A}(\mathsf{x}, r, k, W)$ is 1. Symmetrically, for every k such that $k \log k < \mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W)$, the value of $\mathsf{A}(\mathsf{x}, r, k, W)$ is 0. Therefore, the value k_0 is such that:

$$k_0 - 1 \leq \mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) \text{ and}$$
$$\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W) \leq k_0 \log k_0.$$

As every minimal covering must contain at least one vertex or edge (depending on whether $\mathsf{x} = \mathsf{v}$ or $\mathsf{x} = \mathsf{e}$) of each model of a maximal packing $\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) \leq \mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W)$, we have the following two equations:

$$\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) \leq k_0 \log k_0 \leq (\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) + 1) \log(\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) + 1) \quad (1)$$
$$\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W) \leq k_0 \log k_0 \leq (\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W) + 1) \log(\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W) + 1). \quad (2)$$

Dividing (1) by $\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W)$ and (2) by $\mathsf{cover}_{\Theta_r}(W)$, we get:

$$1 \leq \frac{k_0 \log k_0}{\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W)} \leq \log(\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W) + 1) + \frac{\log \mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W)}{\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W)} = O(\log(\mathsf{x}\text{-}\mathsf{pack}_{\Theta_r}(W))) \text{ and}$$
$$1 \leq \frac{k_0 \log k_0}{\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W)} \leq \log(\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W) + 1) + \frac{\log \mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W)}{\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W)} = O(\log(\mathsf{x}\text{-}\mathsf{cover}_{\Theta_r}(W))).$$

Therefore the value $k_0 \log k_0$ is both an $O(\log OPT)$-approximation of $\mathsf{x\text{-}pack}_{\Theta_r}(W)$ and $\mathsf{cover}_{\Theta_r}(W)$. The value k_0 can be found by performing a binary search in the interval $[\![1, n]\!]$, with $O(\log n)$ calls to Algorithm A. Hence, our approximation algorithm runs in $O(n \cdot \log(n) \cdot m)$ steps. □

5 Erdős–Pósa property and tree-partition width

Using the machinery introduced in Section 3, we are able to prove Theorem 3.

Proof of Theorem 3. Let $k = \mathsf{x\text{-}pack}_{\mathcal{H}}(G)$. This proof is similar to the one of Theorem 5 (progress or reduce). We start with $P := \emptyset$ and $G_0 := G$ and repeatedly apply the Algorithm of Lemma 1 to G_i to either obtain a smaller graph G_{i+1} with the same parameters $\mathsf{x\text{-}cover}_{\mathcal{H}}$ and $\mathsf{x\text{-}pack}_{\mathcal{H}}$ (reduce), or find an $\mathcal{H}$-model M with at most $f_2(h,t)$ edges, in which case we set $G_{i+1} := G_i \setminus A_{\mathsf{x}}(M)$ and $P := P \cup \{M\}$ (progress) and continue. We stop when the current graph has at most $f_2(h,t)$ vertices. Let G_j be this graph. In the end, P contains at most k $\mathcal{H}$-models. Therefore, $C := \bigcup A_{\mathsf{x}}(P) \cup A_{\mathsf{x}}(G_j)$ is an x-$\mathcal{H}$-covering of G of size $(k+1) \cdot f_2(h,t)$, as required. □

The following corollary can be obtained by setting $\mathcal{H} = \mathsf{ex}(H)$.

Corollary 5. *For every graph H and every $t \in \mathbb{N}$, the class of H-minor-models has the (edge and vertex) Erdős–Pósa property for graphs of tree-partition width at most t. Furthermore the gap is a linear function.*

We define $\Theta_{r,r'} = \mathsf{ex}(\theta_{r,r'})$. The rest of this section is devoted to the proof of Theorem 4. Prior to this, we need to introduce a result of Ding *et al.* [DO96].

Tree-partition width has been studied in [See85, Hal91, DO96]. In particular, the authors of [DO96] characterized the classes of graphs of bounded tree-partition width in terms of excluded topological minors. The statement of this result requires additional definitions.

Walls, fans, paths, and stars. The n-wall is the graph with vertex set $[\![1, n]\!]^2$ and whose edge set is:

$$\begin{aligned} &\{\{(i,j),(i,j+1)\},\ 1 \le i,j \le n\} \\ \cup\, &\{\{(2i-1,2j+1),(2i,2j+1)\},\ 1 < 2i \le n \text{ and } 1 \le 2j+1 \le n\} \\ \cup\, &\{\{(2i,2j),(2i+1,2j)\},\ 1 \le 2i < n \text{ and } 1 \le 2j \le n\}\,. \end{aligned}$$

The 7-wall is depicted in Figure 6. The *n-fan* is the graph obtained by adding a dominating vertex to a path on n vertices. A collection of paths is said to be *independent* if two paths of the collection never share interior vertices. The *n-star* is the graph obtained by replacing every edge of $K_{1,n}$ with n independent paths of two edges. The *n-path* is the graph obtained by replacing every edge of an n-edge path with n independent paths

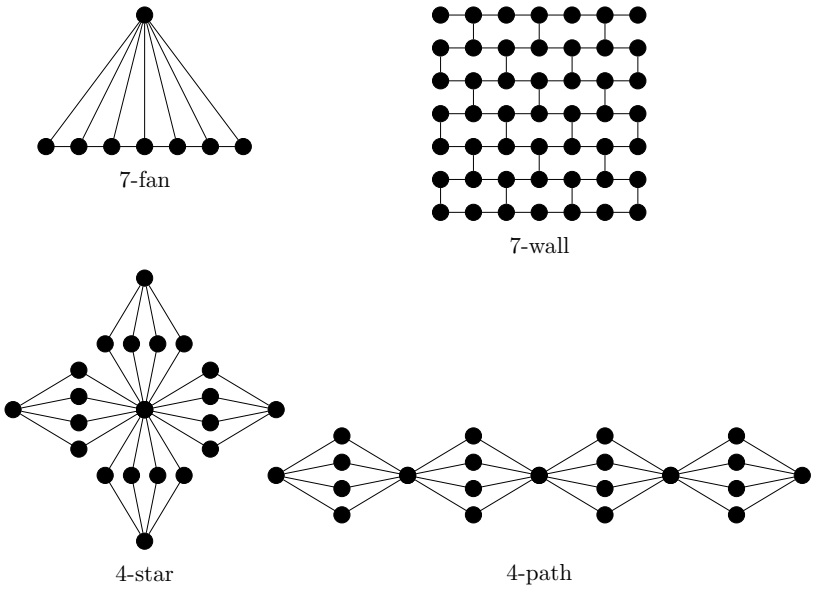

Figure 6: Unavoidable patterns of graphs of large tree-partition width.

of two edges. Examples of these graphs are depicted in Figure 6. The *wall number* (resp. *fan number*, *star number*, and *path number*) of a graph G is defined as the largest integer k such that G contains a model of a k-wall (resp. of a k-fan, of a k-star, of a k-path), or infinity is no such integer exists. Let $\gamma(G)$ denote the maximum of the wall number, fan number, star number, and path number of a graph G.

We need the following result.

Proposition 3 ([DO96])**.** *There is a function $f_{10}\colon \mathbb{N} \to \mathbb{N}$ such that every graph G satisfies $\mathbf{tpw}(G) \leq f_{10}(\gamma(G))$.*

In other words, for every integer k, every graph of large enough tree-partition width contains a model of one of the following graphs: the k-wall, the k-fan, the k-path, or the k-star.

Notice that for every $r, r' \in \mathbb{N}$, $r' \leq r$, the graph $\theta_{r,r'}$ is a minor of the following graphs: the r-path, the r-star, the $(r + r' + 1)$-fan, and the r-wall (for $r \geq 6$). Hence, every graph of large enough tree-partition width contains a $\Theta_{r,r'}$-model. This can easily be generalized to edge-disjoint packings, as follows.

Lemma 11. *For every $r, r' \in \mathbb{N}$, $r' \leq r$, and every $k \in \mathbb{N}_{\geq 1}$, every graph G satisfying $\gamma(G) \geq k(r + r' + 2) - 1$ contains an e-$\Theta_{r,r'}$-packing of size k.*

Using Proposition 3, we get the following corollary.

Corollary 6. *For every $r, r' \in \mathbb{N}$, $r' \leq r$, and every $k \in \mathbb{N}_{\geq 1}$, every graph G satisfying $\mathbf{tpw}(G) \geq f_{10}(k(r + r' + 2) - 1)$ contains an e-$\Theta_{r,r'}$-packing of size k.*

We are now able to give the proof of Theorem 4.

Proof of Theorem 4. According to Corollary 6, for every $k \in \mathbb{N}$, there is a number t_k such that every graph G with $\mathsf{e\text{-}cover}_{\Theta_{r,r'}}(G) = k$ satisfies $\mathbf{tpw}(G) \leq t_k$. Indeed, such a graph does not contain a packing of $k+1$ $\Theta_{r,r'}$-models. Then by Theorem 3 the value $\mathsf{e\text{-}cover}_{\Theta_{r,r'}}(G)$ is bounded above by $f_2(h, t_k) \cdot \mathsf{e\text{-}pack}_{\Theta_{r,r'}}(G)$, and this concludes the proof. □

6 Concluding remarks

The main algorithmic contribution of this paper is a $\log(\mathsf{OPT})$-approximation algorithm for the parameters $\mathsf{v\text{-}pack}_{\theta_r}$, $\mathsf{v\text{-}cover}_{\theta_r}$, $\mathsf{e\text{-}pack}_{\theta_r}$, and $\mathsf{e\text{-}cover}_{\theta_r}$, for every positive integer r. This improves the results of [JPS+14] in the case of vertex packings and coverings and is the first approximation algorithm for the parameters $\mathsf{e\text{-}pack}_{\theta_r}$ and $\mathsf{e\text{-}cover}_{\theta_r}$ for general r. Our proof uses a reduction technique of independent interest, which is not specific to the graph θ_r and can be used for any other (classes of) graphs.

On the combinatorial side, we optimally improved the gap of the edge-Erdős-Pósa property of minor models of θ_r for every r. Also, we were able to show that every class of graphs has the (edge and vertex) Erdős-Pósa property in graphs of bounded tree-partition width, with linear gap. An other outcome of this work is that minor models of $\theta_{r,r'}$ have the edge-Erdős-Pósa. Recall that prior to this work, the only graphs for which this was known were θ_r's.

As mentioned in [RST13], the planarity of a graph H is a necessary condition for the minor models of H to have the edge-Erdős-Pósa property. However, little is known on which planar graphs have this property and with which gap. This is the first direction of research that we want to highlight here. Also, the question of an approximation algorithm can be asked for packing and covering the minor models of different graphs. It was proved in [CC13] that the gap of the vertex-Erdős-Pósa property of minor models of every planar graph is $O(k \operatorname{polylog} k)$. It would be interesting to check if these results can be used to derive a $\operatorname{polylog}(\mathsf{OPT})$-approximation for vertex packing and covering minor models of any planar graph.

Notice that all our results are strongly exploiting Lemma 1 that holds for every finite collection $\mathcal{H}$ of connected graphs. Actually, what is missing in order to have an overall generalization of all of our results, is an extension of Proposition 1 where Θ_r is replaced by any finite collection $\mathcal{H}$ of connected planar graphs. This is an an interesting combinatorial problem even for particular instantiations of $\mathcal{H}$.

References

[ACR03] Alessandro Panconesi Alberto Caprara and Romeo Rizzi. Packing cycles in undirected graphs. *Journal of Algorithms*, 48(1):239 – 256, 2003.

[BBR07] Etienne Birmelé, J. Adrian Bondy, and Bruce A. Reed. The Erdős–Pósa property for long circuits. *Combinatorica*, 27:135–145, 2007.

[BFL+09a] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) kernelization. In *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 629–638, Washington, DC, USA, 2009. IEEE Computer Society.

[BFL+09b] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) kernelization. *CoRR*, abs/0904.0727, 2009.

[BTV07] Cristina Bazgan, Zsolt Tuza, and Daniel Vanderpooten. Efficient algorithms for decomposing graphs under degree constraints. *Discrete Applied Mathematics*, 155(8):979 – 988, 2007.

[BTY11] Hans L. Bodlaender, Stéphan Thomassé, and Anders Yeo. Kernel bounds for disjoint cycles and disjoint paths. *Theoretical Computer Science*, 412(35):4570–4578, 2011.

[CC13] Chandra Chekuri and Julia Chuzhoy. Large-treewidth graph decompositions and applications. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing (STOC)*, pages 291–300, New York, NY, USA, 2013. ACM.

[CRST15] Dimitris. Chatzidimitriou, Jean-Florent Raymond, Ignssi Sau, and Dimitrios M. Thilikos. Minors in graphs of large θ_r-girth. *CoRR*, abs/1510.03041, 2015.

[Die05] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, third edition, 2005.

[DiKW12] Reinhard Diestel, Ken ichi Kawarabayashi, and Paul Wollan. The Erdős–Pósa property for clique minors in highly connected graphs. *Journal of Combinatorial Theory, Series B*, 102(2):454 – 469, 2012.

[DO96] Guoli Ding and Bogdan Oporowski. On tree-partitions of graphs. *Discrete Mathematics*, 149(1–3):45 – 58, 1996.

[EP65] Paul Erdős and Louis Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965.

[FJP10] Samuel Fiorini, Gwenaël Joret, and Ugo Pietropaoli. Hitting diamonds and growing cacti. In *Integer Programming and Combinatorial Optimization*,

volume 6080 of *Lecture Notes in Computer Science*, pages 191–204. Springer Berlin Heidelberg, 2010.

[FJS13] Samuel Fiorini, Gwenaël Joret, and Ignasi Sau. Optimal Erdős–Pósa property for pumpkins. Manuscript, 2013.

[FJW13] Samuel Fiorini, Gwenaël Joret, and David R. Wood. Excluded forest minors and the erdős-pósa property. *Combinatorics, Probability & Computing*, 22(5):700–721, 2013.

[FLMS12] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar F-deletion: Approximation, kernelization and optimal FPT-algorithms. In *IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 470–479, Oct 2012.

[GK09] Jim Geelen and Kasper Kabell. The Erdős–Pósa property for matroid circuits. *Journal of Combinatorial Theory, Series B*, 99(2):407–419, March 2009.

[Gol15] Petr A. Golovach. Personal communication, 2015.

[Hal91] R. Halin. Tree-partitions of infinite graphs. *Discrete Mathematics*, 97(1–3):203 – 217, 1991.

[JPS+14] Gwenaël Joret, Christophe Paul, Ignasi Sau, Saket Saurabh, and Stéphan Thomassé. Hitting and harvesting pumpkins. *SIAM Journal on Discrete Mathemarics*, 28(3):1363–1390, 2014.

[KK12] Ken-Ichi Kawarabayashi and Yusuke Kobayashi. Edge-disjoint odd cycles in 4-edge-connected graphs. In *29th International Symposium on Theoretical Aspects of Computer Science, (STACS), February 29th - March 3rd, 2012, Paris, France*, pages 206–217, 2012.

[KKK12] Naonori Kakimura, Ken-ichi Kawarabayashi, and Yusuke Kobayashi. Erdős-pósa property and its algorithmic applications: Parity constraints, subset feedback set, and subset packing. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1726–1736. SIAM, 2012.

[KLL02] T. Kloks, C.M. Lee, and Jiping Liu. New algorithms for k-face cover, k-feedback vertex set, and k-disjoint cycles on plane and planar graphs. In *Graph-Theoretic Concepts in Computer Science*, volume 2573 of *Lecture Notes in Computer Science*, pages 282–295. Springer Berlin Heidelberg, 2002.

[KN07] Ken-Ichi Kawarabayashi and Atsuhiro Nakamoto. The erdős–pósa property for vertex- and edge-disjoint odd cycles in graphs on orientable surfaces. *Discrete Mathematics*, 307(6):764 – 768, 2007.

[KNS+07] Michael Krivelevich, Zeev Nutov, Mohammad R. Salavatipour, Jacques V. Yuster, and Raphael Yuster. Approximation algorithms and hardness results for cycle packing problems. *ACM Transactions on Algorithms*, 3(4), November 2007.

[KNY05] Michael Krivelevich, Zeev Nutov, and Raphael Yuster. Approximation algorithms for cycle packing problems. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 556–561, Philadelphia, PA, USA, 2005. SIAM.

[KV04] Daniel Král' and Heinz-Jürgen Voss. Edge-disjoint odd cycles in planar graphs. *Journal of Combinatorial Theory, Series B*, 90(1):107 – 120, 2004.

[PW12] M. Pontecorvi and Paul Wollan. Disjoint cycles intersecting a set of vertices. *Journal of Combinatorial Theory, Series B*, 102(5):1134 – 1141, 2012.

[RR01] Dieter Rautenbach and Bruce Reed. The Erdős–Pósa property for odd cycles in highly connected graphs. *Combinatorica*, 21:267–278, 2001.

[RRST96] Bruce A. Reed, Neil Robertson, Paul D. Seymour, and Robin Thomas. Packing directed circuits. *Combinatorica*, 16(4):535–554, 1996.

[RS86] Neil Robertson and Paul D. Seymour. Graph Minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(2):92–114, 1986.

[RST13] Jean-Florent Raymond, Ignasi Sau, and Dimitrios M. Thilikos. An edge variant of the Erdős–pósa property. *CoRR*, abs/1311.1108, 2013.

[See85] Detlef Seese. Tree-partite graphs and the complexity of algorithms. In Lothar Budach, editor, *Fundamentals of Computation Theory*, volume 199 of *Lecture Notes in Computer Science*, pages 412–421. Springer Berlin Heidelberg, 1985.

[SV05] Mohammad R. Salavatipour and Jacques Verstraete. Disjoint cycles: Integrality gap, hardness, and approximation. In *Integer Programming and Combinatorial Optimization*, volume 3509 of *Lecture Notes in Computer Science*, pages 51–65. Springer Berlin Heidelberg, 2005.